

Zakk Lefkowitz

SENIOR SOFTWARE ENGINEER

📍 Richmond, VA (EST) ✉ zlefkowitz@gmail.com 🌐 zakk.io/resume 📄 in/zakklefkowitz 🔄 zakk113

Summary

Senior Software Engineer at **Amazon Fire TV**. 9 years in distributed systems, mobile-cloud architecture, and data engineering. Launched the Discovery Catalog data platform; previously shipped Continue Watching to **100M+ devices** as client-side tech lead. Consistent record of championing the harder, better architectural decision to build robust systems at Amazon scale. Pushing agentic engineering practice at Fire TV; built and operate a self-iterating agent orchestrator shipping **~3x** prior-year output of peer-reviewed code.

Experience

Amazon, Fire TV

SENIOR SOFTWARE ENGINEER

Richmond, VA

April 2020 - Present

Tech lead for Catalog Enrichment, previously client-side tech lead on Partner Integrations. Promoted to Sr. Software Engineer (Q4 2022).

Discovery Catalog Platform

Data Engineering, Distributed Systems

- Designed and built a governed data platform extending the Fire TV Catalog beyond streaming-provider feeds to enriched sources (Wikipedia articles, scene descriptions), replacing fragmented ad-hoc pulls by downstream teams.
- Drove the architectural decision to publish to a separate output stream rather than extending the core catalog, unlocking a **10x data-volume ceiling** and a faster schema-evolution path for novel data types.
- Aligned the Fire TV Search team to build its Prime Video scene-data curation app on the new platform, setting precedent for centralized data governance across downstream teams.
- Delivered to production solo in three months against a documented 4.5-month estimate. Enabled by agentic engineering.

LLM Genre Classification

LLM Applications & Evaluation, AWS Bedrock, Data Engineering

- Built an LLM genre classifier that replaced unreliable provider data, embedded in the catalog's streaming pipeline at **~\$0.001 per entity**; engineered an eligibility filter scoping inference to the ~5% of updates, keeping long-tail content off the latency-sensitive path.
- Invented **CMAR (Catalog Metadata Acceptance Rate)**: a measurement methodology using an LLM judge calibrated against PM ratings (>90% agreement on a 10K-entity sample) to validate enrichment quality. Classifier hit **95% CMAR vs. 82% baseline**; the metric earned PM confidence and unlocked downstream customer-team adoption of LLM-classified genre values.
- Tech lead on subsequent enrichment features shipped by mid-level engineers; core contributor to the org's shared Bedrock integration library.

Third-Party Continue Watching on Fire TV

Mobile-Cloud Architecture, Rust, Android, Java Services

- Led the client-side architecture on the VP-level initiative that gave Fire TV visibility into what customers watched inside partner apps, enabling Continue Watching and **8-figure incremental ad revenue**.
- Drove the decision to build the on-device middleware in Rust rather than piggyback on existing Android services. Shipped to 100M+ devices at <5MB memory footprint, proving out the native-service pattern subsequently adopted across Fire TV.
- Drove the decision to build a custom Amazon SDK rather than adopt Android's standard Watch Next API. Owned the entire partner-facing surface end-to-end (SDK, [integration docs](#), [reference sample app](#)); the resulting API surface became a partner-cited deal-closing factor across **Disney+, HBO Max, Hulu, Starz, and MGM+**.
- Built and operated the cloud ingestion backend handling sensitive customer and partner data at **~200K TPS peak**.
- Owned the privacy/consent surface end-to-end; copy experimentation drove **opt-in from 30% → 75%** in months.

Agentic Engineering

Agentic Workflow Automation, MCP

- Authored and maintain **catalog-engineer**: a shared agent spec (context + MCP servers + skills) adopted 100% across a 3-team, ~40-person org, cited by VP as a best-practice example, and replicated by sister Fire TV orgs.
- Built and operate an internal **human-gated autonomous engineering pipeline**. Shipping production code through Amazon's standard peer code review. **Q1 2026 alone shipped ~70% as many merged PRs as all of 2025** (~3x output multiplier). Shared the system in an org-wide talk to 100+ engineers, driving 10+ direct adopters and two VP follow-ups.
- Trusted by Fire TV VP leadership to help provision a new AI Engineering velocity team; subsequently drove Catalog org adoption of a code package context graph system, landing **100 core packages** in 30 days.

CapTech Consulting

SENIOR CONSULTANT

Richmond, VA

July 2017 - March 2020

Led cloud modernization initiatives for Global Top 10 insurance client. Promoted to Senior Consultant (Q3 2019).

Insurance Cloud Modernization

AWS ECS, CodeDeploy, Terraform, TypeScript, Angular, Node.js

- Architected **AWS ECS** migration and **CodeDeploy**-based CI/CD for a legacy Windows-based Java payment application processing **~\$1bn/year**; built supporting **Terraform** IaC and a TypeScript deployment CLI.

Skills

Languages & Technologies
AWS, Cloud & Infrastructure
Architecture & Core
Competencies

Rust, Java, Kotlin, TypeScript, Python, SQL, Android, Spring Boot, Apache Spark, RocksDB, MCP
EC2, ECS, Lambda, DynamoDB, SQS, Bedrock, CDK, S3, CloudWatch, CodeDeploy, Terraform (IaC)
Distributed Systems, Data Engineering, Mobile-Cloud Architecture, Agentic Workflow Automation,
LLM Applications & Evaluation

Education

Virginia Polytechnic Institute and State University

B.S. COMPUTER SCIENCE, MATHEMATICS MINOR, CUM LAUDE

Blacksburg, VA

May 2017